


[Sign in](#)

Web Results 1 - 10 of about 105 for reinforcement learning garbage collection sarsa algorithm. (0.54 sec)

Scholarly articles for reinforcement learning garbage collection sarsa algorithm



Online Resource Allocation Using Decompositional ... - by Tesauro - 1 citations
 To Collect or Not to Collect? Machine Learning for ... - by Andreasson - 1 citations
 Solving Complex Problems with Reinforcement Learning - by Randløv - 4 citations

[PDF] Decompositional Reinforcement Learning and Workload Management

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Keywords: **reinforcement learning**, distributed AI, autonomic computing. Abstract ... facts (eg Java **garbage collection**). Encoding of the state ...

www.cs.utexas.edu/~nalini/research/tesauro-decomp.pdf - [Similar pages](#)

[PDF] To Collect or Not To Collect? Machine Learning for Memory Management

File Format: PDF/Adobe Acrobat - [View as HTML](#)

two **algorithms** are well suited for the **garbage collection** problem. ...
reinforcement learning might improve existing **garbage collection** techniques. ...

www.nada.kth.se/~hoffmann/jvm02.pdf - [Similar pages](#)

JVM '02 paper

Reinforcement learning provides an approach in which an agent interacts with ...

Both the other two **algorithms** are well suited for the **garbage collection** ...

www.usenix.org/events/javavm02/full_papers/andreasson/andreasson_html/ - 60k - [Cached](#) - [Similar pages](#)

Bibliography

... Menezes and Alan Wood, Distributed **Garbage Collection** of Tuple ... A very good (formal) introduction to **reinforcement learning**. ... **SARSA** and TD-learning are explained ...

bio6.itek.norut.no/werner/Phd/PhdOnTheWeb/node22.html - 39k - Supplemental Result - [Cached](#) - [Similar pages](#)

Bibliography

MW99: Ronaldo Menezes and Alan Wood, Distributed **Garbage Collection** of Tuple ...

SAG98: Richard S. Sutton and Andrew G. Barto, **Reinforcement Learning**: An ...

bio6.itek.norut.no/werner/phd03/PhdOnTheWeb/node22.html - 39k - [Cached](#) - [Similar pages](#)

[PDF] Solving Complex Problems with Reinforcement Learning

File Format: PDF/Adobe Acrobat - [View as HTML](#)

6 Combining **Reinforcement Learning** with a Local Control **Algorithm** 60 ... 1995)
 used the **Sarsa** and **Sarsa(λ) algorithm** (called modified Q-learning in his ...

www.nbi.dk/~randlov/thesis.pdf - [Similar pages](#)

Colloquium Papers - Artificial Intelligence @ RightNow Technologies

We then discuss an on-line **learning algorithm**, based on the "Hedge" ...
 statistical "**garbage collection**" for new event detection, clustering in time to ...


ai.rightnow.com/colloquium/papers.php - 273k - [Cached](#) - [Similar pages](#)

[PDF] Cooperation through Reinforcement Learning

File Format: PDF/Adobe Acrobat - [View as HTML](#)

that it is possible To implement a **Reinforcement Learning algorithm** in continuous ...
 ... here we present an **algorithm** that is noticeably similar To **SARSA**. ...

<http://www.google.com/search?hl=en&lr=&q=reinforcement+learning+garbage+collection+s...> 1/25/06

[Sign in](#)

Web Results **81 - 87** of about **105** for **reinforcement learning garbage collection sarsa algorithm**. (0.15 sec)

consumed.5gigs.com/files/words.txt

2922k - Supplemental Result - [Cached](#) - [Similar pages](#)

hanzabu.dyndns.org/pub/NetBSD/arch/hpcmips/snapsho...

2430k - Supplemental Result - [Cached](#) - [Similar pages](#)

[PDF] [Recent Advances in Hierarchical Reinforcement Learning](#)

File Format: PDF/Adobe Acrobat -

Reinforcement Learning. DP algorithms have complexity polynomial in the ...

Both the Q-learning and Sarsa learning algorithms also apply to SMDPs, both ...

www.ingentaconnect.com/content/klu/disc/2003/00000013/F0020001/05110821 - [Similar pages](#)

[PDF] [Recent Advances in Hierarchical Reinforcement Learning](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Both the Q-learning and Sarsa learning algorithms also apply to SMDPs, both

continuous- ... Average reward reinforcement learning: Foundations, algorithms, ...

www.cs.utexas.edu/users/kuipers/readings/Barto-hrl-03.pdf - [Similar pages](#)

[[More results from www.cs.utexas.edu](#)]

[PDF] www.lib.utexas.edu/etd/d/2003/gomezjf036/gomezjf03...

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Supplemental Result - [Similar pages](#)

[PDF] [HIERARCHICAL REINFORCEMENT LEARNING IN CONTINUOUS STATE AND MULTI ...](#)

File Format: PDF/Adobe Acrobat - [View as HTML](#)

6.3 A Hierarchical Multi-Agent Reinforcement Learning Algorithm Online

algorithms such as SARSA and Q-learning also generalize to the SMDP case ...

www-anw.cs.umass.edu/pubs/2005/ghavamzadeh_thesis05.pdf - [Similar pages](#)

[PDF] <ftp://ftp.cs.utexas.edu/pub/AI-Lab/tech-reports/UT...>

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Supplemental Result - [Similar pages](#)

In order to show you the most relevant results, we have omitted some entries very similar to the 87 already displayed.

If you like, you can repeat the search with the omitted results included.

◀ Goooooooooogle


Result Page: [Previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#)

reinforcement learning garbage colle

[Search within results](#) | [Language Tools](#) | [Search Tips](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)

Web Results 1 - 10 of about 217 for **java virtual machine sarsa reinforcement learning**. (0.57 seconds)

Java Virtual Machine 2002 Symposium, San Francisco, CA

Eva explained how **Reinforcement Learning** contributes to **Java Virtual Machine** autonomous decision making to perform garbage collection. ...

home.comcast.net/~osorioj/JVM02Rev.html - 33k - [Cached](#) - [Similar pages](#)

[PDF] To Collect or Not To Collect? Machine Learning for Memory Management

File Format: PDF/Adobe Acrobat - [View as HTML](#)

and requirements of a dynamic **Java Virtual Machine**. ... **Reinforcement Learning** for a. dynamic JVM. KTH/Appeal **Virtual Machines**,. Stockholm, Sweden. ...

www.nada.kth.se/~hoffmann/jvm02.pdf - [Similar pages](#)

[PDF] inside:

File Format: PDF/Adobe Acrobat

IBM Centre for **Java** Technology. "**Virtual Machine** technology is here to ...

the **SARSA** approach yields an optimal. **learning** scheme for adaptive garbage ...

www.usenix.org/publications/ library/proceedings/javavm02/JavaVM2002.pdf - [Similar pages](#)

JVM '02 paper

JRockit™, the **Java™ Virtual Machine** (JVM) constructed by Appeal **Virtual Machines**

... **Reinforcement Learning** for a dynamic JVM. KTH/Appeal **Virtual Machines**, ...

www.usenix.org/events/javavm02/ full_papers/andreasson/andreasson_html/ - 60k -

[Cached](#) - [Similar pages](#)

International Conference on Computational Intelligence for ...

Virtual Machine Code for Embedded Systems SL .Yun, DK. Nam, JS Kim, SM. Oh ...

Fuzzy **Sarsa**: An approach to fuzzifying **Sarsa Learning** L. Tokarchuk ...

www.ise.canberra.edu.au/conferences/ cimca04/Conference_Prog.htm - 31k - [Cached](#) - [Similar pages](#)

Bibliography

... A case for grid computing on **virtual machines**, Proceedings International ... is reinvented with the start of **Java**. ... for Genetic Algorithms, **Machine Learning** 9 (1992 ...

bio6.itek.norut.no/werner/Phd/PhdOnTheWeb/node22.html - 39k - Supplemental Result - [Cached](#) - [Similar pages](#)

Bibliography

... P. Dinda, and J. Fortes, A case for grid computing on **virtual machines**, ...

SAG98: Richard S. Sutton and Andrew G.Barto, **Reinforcement Learning**: An ...

bio6.itek.norut.no/werner/phd03/PhdOnTheWeb/node22.html - 39k - [Cached](#) - [Similar pages](#)

[PDF] Reinforcement Learning Benchmarks and Bake-offs II

File Format: PDF/Adobe Acrobat - [View as HTML](#)

Optimal **Reinforcement Learning**", Journal of **Machine Learning** Research 3 (2002)

213-231. ... platforms supporting the **Java Virtual Machine**. ...

www.cs.rutgers.edu/~mlittman/ topics/nips05-mdp/bakeoffs05.pdf - [Similar pages](#)

[PDF] REINFORCEMENT LEARNING OF STRATEGIES FOR SETTLERS OF CATAN

File Format: PDF/Adobe Acrobat - [View as HTML](#)

In this paper we study the application of **machine. learning** methods in complex computer games. A. combination of hierarchical **reinforcement learning** and ...

eprints.pascal-network.org/archive/ 00000425/01/LAG-37_pfeiffer_2004.pdf - [Similar pages](#)

Karan's Webpages - Projects

Present a performance comparison between two popular **Reinforcement Learning** algorithms: the **Sarsa**-Lambda algorithm and Watkin's Q-Lambda algorithm. ...

www.cs.ttu.edu/~gupta/Computers/ - 64k - [Cached](#) - [Similar pages](#)

Try your search again on [Google Book Search](#)

Info when you want it, right on your desktop

Free! [Download Google Desktop](#)

Goooooooooooooogle ►

Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied?](#) [Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)[Web](#) [Images](#) [Groups](#) [News](#) [Froogle](#) [Local](#) [more »](#)

java virtual machine run time

[Advanced Search](#)
[Preferences](#)**Web**Results 1 - 10 of about 5,720,000 for **java virtual machine run time**. (0.11 seconds)**Downloads Home**

Also includes **Java** Development Kit and **Java Runtime** Environment. ... **Java** Foundation Class (JFC) / Swing, **Java** HotSpot **Virtual Machine**, **Java** Media Framework ...
java.sun.com/downloads/ - 34k - Jan 23, 2006 - [Cached](#) - [Similar pages](#)

Java 2 Platform, Standard Edition (J2SE)

... and **running** applets and applications in the **Java** programming language. ...
Java SE is also at home in today's more demanding embedded and real-time ...
java.sun.com/j2se/ - 32k - Jan 23, 2006 - [Cached](#) - [Similar pages](#)
[[More results from java.sun.com](#)]

Download Java software from Sun Microsystems

Get the latest version of the Sun Microsystems **Java Runtime** Environment (JRE)
... To complete the download of the **Java Runtime** Environment, please select ...
www.java.com/en/download/manual.jsp - 19k - Jan 23, 2006 - [Cached](#) - [Similar pages](#)

Microsoft Java Virtual Machine Support

... to provide a smooth transition for current users of the Microsoft **Java Virtual Machine** (MSJVM), ... **Java Runtime** Solution, Description. Sun Microsystems ...
www.microsoft.com/mscorp/java/ - 22k - [Cached](#) - [Similar pages](#)

Apple - Mac OS X - Java

Recognizing this, Apple has made **Java** a core component of Mac OS X. Mac OS X ...
you have the **Java** Developer Kit (JDK) and the HotSpot **virtual machine** (VM) ...
www.apple.com/macosx/features/java/ - 26k - Jan 23, 2006 - [Cached](#) - [Similar pages](#)

Java Virtual Machine - All about JVMs

There are four major sections: **Java** development kits and **Java runtime** environments, standalone **Java virtual machines**, **Java** products for embedded systems and ...
java-virtual-machine.net/ - 4k - [Cached](#) - [Similar pages](#)

Java virtual machines, Java development kits and Java runtime ...

This page is supposed to give an overview of **Java virtual machines** (JVMs), **Java** development kits (JDKs), **Java runtime** environments (JREs) and related ...
java-virtual-machine.net/other.html - 37k - [Cached](#) - [Similar pages](#)
[[More results from java-virtual-machine.net](#)]

Download Java Software

Download **Java** Software for your desktop computer now! **Java Runtime** Environment Version 5.0 Update 6 Download **Java** Software Now ...
java.com/getjava/index.jsp - 10k - Jan 23, 2006 - [Cached](#) - [Similar pages](#)
[[More results from java.com](#)]

Kaffe.org

Kaffe is a clean room implementation of the **Java virtual machine**, plus the associated class libraries needed to provide a **Java runtime** environment. ...
www.kaffe.org/ - 9k - [Cached](#) - [Similar pages](#)

List of Java virtual machines, Java development kits and Java ...

This page gives an overview of **Java virtual machines** (JVM), **Java development kits** (JDK), **Java runtime** environments (JRE) and related products. ...

schmidt.devlib.org/java/virtual-machines.html - 6k - [Cached](#) - [Similar pages](#)

Try your search again on [Google Book Search](#)

Goooooooooooooogle ►

Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

Free! Get the Google Toolbar. [Download Now](#) - [About Toolbar](#)




java virtual machine run time Search

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

[Sign in](#)



Web Results 1 - 10 of about 10,400 for **2001 virtual machine memory reinforcement learning garbage colle**

Scholarly articles for 2001 virtual machine memory reinforcement learning garbage collection process



[To Collect or Not to Collect? Machine Learning for ...](#) - by Andreasson - 1 citations

[Using Code Collection to Support Large Applications on ...](#) - by Popa - 2 citations

[Machine Learning for Memory Management](#) - by Andreasson - 0 citations

Citations: Parallel garbage collection for shared memory ...

Parallel **garbage collection** for shared **memory** multiprocessors. In Proceedings of the Java **Virtual Machine** Research and Technology Symposium, Monterey, ...

citeseer.ist.psu.edu/context/1918766/489697 - 13k - [Cached](#) - [Similar pages](#)

Citations: Hot-swapping between a mark&sweep and a mark&compact ...

... of the Java **Virtual Machine** Research and Technology Symposium, 2001. ...

4 Reinforcement Learning Reinforcement learning methods solve a class of ...

citeseer.ist.psu.edu/context/1734791/0 - 12k - [Cached](#) - [Similar pages](#)

[PDF] To Collect or Not To Collect? Machine Learning for Memory Management

File Format: PDF/Adobe Acrobat - [View as HTML](#)

garbage collection always intrudes on the **process** time ... **Reinforcement Learning** for a. dynamic JVM. KTH/Appeal **Virtual Machines**,. Stockholm, Sweden. ...

www.nada.kth.se/~hoffmann/jvm02.pdf - [Similar pages](#)

Java Opportunities and Challenges in a 64-bit World: Solutions ...

It uses 64 bits to address **virtual memory**, allowing each **process** ... in research to combine this unified **collector** with a **reinforcement learning** approach to ...

www.devx.com/Intel/Article/17457 - 40k - [Cached](#) - [Similar pages](#)

Java Opportunities and Challenges in a 64-bit World: Solutions ...

... with some interesting challenges—requires a clever Java **Virtual Machine**. ...

in research to combine this unified **collector** with a **reinforcement learning** ...

www.devx.com/Intel/Article/17457/1763 - 30k - [Cached](#) - [Similar pages](#)

JVM '02 paper

The ratio of freed **memory** after completed **garbage collection** compared to the ...

Reinforcement Learning for a dynamic JVM. KTH/Appeal **Virtual Machines**, ...

www.usenix.org/events/javavm02/full_papers/andreasson/andreasson_html/ - 60k - [Cached](#) - [Similar pages](#)

[PDF] Using Code Collection to Support Large Applications on Mobile Devices

File Format: PDF/Adobe Acrobat - [View as HTML](#)

The Java **virtual machine** [4] uses classes as a unit for code. management. ...

memory, the **collection process** is used frequently and usually the ...

www.cs.ucl.ac.uk/staff/C.Raiciu/files/p2858-popa.pdf - [Similar pages](#)

[PDF] Using Code Collection to Support Large Applications on Mobile Devices

File Format: PDF/Adobe Acrobat - [View as HTML](#)

memory, the **collection process** is used frequently and usually the ... **Machine**

Learning for Memory. Management, in Proceedings of the 2nd Java **Virtual** ...

www.eecs.berkeley.edu/~popa/p1568932858-popa.pdf - [Similar pages](#)

[ANU - Faculty of Engineering and Information Technology \(FEIT ...](#)

In 18th International Conference on **Machine Learning ICML, 2001**. ... **Garbage collection** is key to modern languages like Java and C#. ...

cs.anu.edu.au/honours/topics.html - 113k - [Cached](#) - [Similar pages](#)

[pocj ACM 2001 Curricula Tree](#)

File Format: Microsoft Word - [View as HTML](#)

Hierarchy of **virtual machines**. 004.002.003. Intermediate languages ... **Reinforcement learning**. 010.009.001. Definition and examples of planning systems ...

www.intag.org/pages/WP/ACM_2001_Curricula.doc - [Similar pages](#)

Try your search again on [Google Book Search](#)

Info when you want it, right on your desktop

Free! [Download Google Desktop](#)

Goooooooooooooogle ►

Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

2001 virtual machine memory reinfo

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied? Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

citeseer.ist.psu.edu/context/49810/0 - 28k - [Cached](#) - [Similar pages](#)

Chapter 6

Describe the **Sarsa algorithm**. Why is this called an on-policy control **algorithm**?

... How does the Q-learning **algorithm** differ from the **Sarsa algorithm**? ...

cs.gettysburg.edu/~tneller/cs391/ch6.html - 6k - [Cached](#) - [Similar pages](#)

Try your search again on [Google Book Search](#)

Info when you want it, right on your desktop

Free! [Download Google Desktop](#)

Goooooooooooooogle ►

Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

sarsa algorithm

Search

[Search within results](#) | [Language Tools](#) | [Search Tips](#) | [Dissatisfied?](#) [Help us improve](#)

[Google Home](#) - [Advertising Programs](#) - [Business Solutions](#) - [About Google](#)

©2006 Google

Recent Advances in Hierarchical Reinforcement Learning (2003) (Make Corrections) (12 citations)

Andrew G. Barto Sridhar Mahadevan Autonomous Learning Laboratory Department...

View or download:

umass.edu/~mahadeva/...dedspaper.pdf
Cached: [PS.gz](#) [PS](#) [PDF](#) [Image](#) [Update](#) [Help](#)

CiteSeer
Electronic Library

[Home/Search](#) [Bookmark](#) [Context](#) [Related](#)

From: umass.edu/~mahad...organizedpubs
(more)
(Enter author homepages)

(Enter summary)

Rate this article: 1 2 3 4 5 (best)
[Comment on this article](#)

Abstract: Reinforcement learning is bedeviled by the curse of dimensionality: the number of parameters to be learned grows exponentially with the size of any compact encoding of a state. Recent attempts to combat the curse of dimensionality have turned to principled ways of exploiting temporal abstraction, where decisions are not required at each step, but rather invoke the execution of temporally-extended activities which follow their own policies until termination. This leads naturally to hierarchical... ([Update](#))

Cited by: [More](#)

Spoken Dialogue Management Using Hierarchical - Reinforcement Learning And ([Correct](#))
Intrinsically Motivated Reinforcement Learning: A.. - Stout, Konidaris, Barto (2005) ([Correct](#))
Hierarchical Reinforcement Learning in.. - Fischer, Rovatsos, Weiss (2004) ([Correct](#))

Active bibliography (related documents): [More](#) [All](#)

7.9: Recent Advances in Hierarchical Reinforcement Learning - Barto, Mahadevan (2003) ([Correct](#))
1.4: Learning and Approximate Dynamic Programming - Scaling.. - Si, Barto, Powell.. ([Correct](#))
1.3: Spatiotemporal Abstraction of Stochastic Sequential Processes - Mahadevan ([Correct](#))

Similar documents based on text: [More](#) [All](#)

0.2: Report of the 1996 Workshop on Reinforcement Learning - Mahadevan, Kaelbling ([Correct](#))
0.2: Extending Hierarchical Reinforcement Learning to.. - Ghavamzadeh, Mahadevan ([Correct](#))
0.2: Hierarchical Multi Agent Reinforcement Learning - Makar, Mahadevan (2001) ([Correct](#))

Related documents from co-citation: [More](#) [All](#)

8: Between mdps and semi-mdps: A framework for temporal abstraction in reinforcemen.. - Sutton, Precup et al. - 1999
6: Reinforcement learning: An introduction - Sutton, Barto - 1998
5: Using collective intelligence to route internet traffic - Wolpert, Tumer et al. - 1999

BibTeX entry: ([Update](#))

Barto, A.G., Mahadevan, S.: Recent advances in hierarchical reinforcement learning. Discrete event systems (2003, to appear) <http://citeseer.ist.psu.edu/barto03recent.html> [More](#)


```
@misc{ barto03recent,  
  author = "A. Barto and S. Mahadevan",  
  title = "Recent advances in hierarchical reinforcement learning",  
  note = "Discrete event systems (2003, to appear)",  
  year = "2003",  
  url = "citeseer.ist.psu.edu/barto03recent.html" }
```

Citations (may not include all citations):

631 Learning from Delayed Rewards (context) - Watkins - 1989
558 Reinforcement Learning: An Introduction - Sutton, Barto - 1998
384 Neuro-Dynamic Programming (context) - Bertsekas, Tsitsiklis - 1996
360 Reinforcement learning: A survey - Kaelbling, Littman et al. - 1996
285 Some studies in machine learning using the game of checkers (context) - Samuel - 1967
285 Some studies in machine learning using the game of checkers (context) - Samuel - 1959
270 Machine Learning (context) - Watkins, Dayan - 1992

- 250 Learning to act using real-time dynamic programming - Barto, Bradtke et al. - 1995
- 219 Dynamic Programming: Deterministic and Stochastic Models (context) - Bertsekas - 1987
- 217 Learning and executing generalized robot plans (context) - Fikes, Hart et al. - 1972
- 182 Markov games as a framework for multi-agent reinforcement le.. - Littman - 1994
- 181 A model for reasoning about persistence and causation (context) - Dean, Kanazawa - 1989
- 172 Planning and acting in partially observable stochastic domai.. - Kaelbling, Littman et al. - 1998
- 134 Introduction to Stochastic Dynamic Programming (context) - Ross - 1983
- 128 Multi-agent reinforcement learning: Independent vs - Tan - 1993
- 108 Generalization in reinforcement learning: Successful example.. - Sutton - 1995
- 102 the convergence of stochastic iterative dynamic programming .. - Jaakkola, Jordan et al. - 1994
- 101 Tractable inference for complex stochastic processes - Boyen, Koller - 1998
- 101 Transition network grammars for natural language analysis (context) - Woods - 1970
- 98 Multiagent Systems: A Modern Approach to Distributed Artific.. (context) - Weiss - 1999
- 80 Reinforcement learning with hierarchies of machines - Parr, Russell - 1997
- 74 Reinforcement Learning with Selective Perception and Hidden.. (context) - McCallum - 1996
- 72 Between mdps and semi-mdps: A framework for temporal abstrac.. - Sutton, Precup et al. - 1999
- 71 Hierarchical reinforcement learning with the maxq value func.. - Dietterich - 2000
- 70 Learning topological maps with weak local odometric informat.. - Shatkay, Kaelbling - 1997
- 59 Toward a modern theory of adaptive networks: Expectation and.. (context) - Sutton, Barto - 1981
- 58 learning using connectionist systems (context) - Rummerly, Niranjan - 1994
- 56 a self-teaching backgammon program (context) - Tesauro - 1994
- 51 Reinforcement learning for dynamic channel allocation in cel.. - Singh, Bertsekas - 1996
- 49 A reinforcement learning method for maximizing undiscounted .. (context) - Schwartz - 1993
- 47 Achieving artificial intelligence through building robots (context) - Brooks - 1986
- 46 Approximate dynamic programming for real-time control and ne.. (context) - Werbos - 1992
- 45 Average reward reinforcement learning: Foundations - Mahadevan - 1996
- 43 Finding structure in reinforcement learning - Thrun, Schwartz - 1994
- 42 The Hierarchical Hidden Markov Model: Analysis and Applicati.. - Fine, Singer et al. - 1998
- 39 Hierarchical Control and Learning for Markov Decision Proces.. - Parr - 1998
- 37 Multi-time models for temporally abstract planning - Precup, Sutton - 1997
- 37 Convergence results for single-step on-policy reinforcement-.. - Singh, Jaakkola et al. - 2000
- 37 The complexity of decentralized control of markov decision p.. - Bernstein, Zilberstein et al. - 2000
- 35 Singular Perturbation Methods in Control: Analysis and Desig.. (context) - Kokotovic, Khalil et al. - 1986
- 35 A heuristic approach to the discovery of macro-operators (context) - Iba - 1989
- 33 Reinforcement learning methods for continuous-time markov de.. - Bradtke, Du - 1994
- 32 Building and understanding adaptive system statisticalnumeri.. (context) - Building, systems et al. - 1987
- 32 A unified framework for hybrid control: Model and optimal co.. - Branicky, Borkar et al. - 1998
- 30 Learning to Solve Problems by Searching for Macro-Operators (context) - Korf - 1985
- 29 Xavier: A robot navigation architecture based on partially o.. (context) - Koenig, Simmons - 1997
- 28 Learning to improve coordinated actions in cooperative distr.. - Sugawara, Lesser - 1998
- 28 A feedback control structure for on-line learning tasks - Huber, Grunpen - 1997
- 24 Reinforcement learning with a hierarchy of abstract models - Singh - 1992
- 24 The Hedonistic Neuron: A Theory of Memory (context) - Klopff - 1982
- 23 Elevator group control using multiple reinforcement learning.. - Crites, Barto - 1998
- 21 Brain function and adaptive systems---A heterostatic theory (context) - Klopff - 1972
- 21 Scaling reinforcement learning toward RoboCup soccer - Stone, Sutton - 2001
- 20 Self-improving factory simulation using continuous-time aver.. - Mahadevan, Marchalleck et al. - 1997
- 17 Theory of Neural-Analog Reinforcement Systems and its Applic.. (context) - Minsky - 1954
- 16 Scaling reinforcement learning algorithms by learning variab.. - Singh - 1992
- 15 Emergent hierarchical control structure Learning reactivehie.. - hierarchical, Learning et al. - 1996
- 15 An analysis of temporal-di#erence learning with function app.. (context) - Tsitsiklis, Van Roy - 1997
- 14 Advanced forecasting methods for global crisis warning and m.. (context) - Werbos - 1977
- 14 Multilayer control of large markov chains (context) - Forestier, Varaiya - 1978
- 14 Theoretical results on reinforcement learning with temporall.. - Precup, Sutton et al. - 1998
- 13 Practical issues in temporal di#erence learning (context) - Tesauro - 1992
- 11 Programmable reinforcement learning agents - Andre, Russell - 2000
- 10 Dynamic Probabilistic Systems: Semi-Markov and Decision Proc.. (context) - Howard - 1971
- 10 Temporal Abstraction in Reinforcement Learning - Precup - 2000
- 10 Large-Scale Dynamic Optimization Using Teams of Reinforcemen.. - Crites - 1996
- 9 Hierarchical multi-agent reinforcement learning - Makar, Mahadevan et al. - 2001
- 9 Learning hierarchical partially observable markov decision p.. - Theocharous, Rohanimanesh et al. - 2001

- 8 Localizing search in reinforcement learning - Grudic, Ungar - 2000
- 7 Automatic discovery of subgoals in reinforcement learning us.. - McGovern, Barto - 2001
- 7 Hierarchical Learning and Planning in Partially Observable M.. (context) - Theodorou - 2002
- 7 Hierarchical memory-based reinforcement learning - Hernandez, Mahadevan - 2001
- 6 Discovering hierarchy in reinforcement learning with hexq - Hengst - 2002
- 6 Markov Decision Problems (context) - Puterman - 1994
- 6 Solving semi-markov decision problems using average reward r.. - Das, Gosavi et al. - 1999
- 5 Autonomous Discovery of Temporal Abstractions from Interacti.. (context) - McGovern - 2002
- 5 Approximate planning with hierarchical partially observable .. - Theodorou, Mahadevan - 2002
- 5 Decision-theoretic planning with concurrent temporally exten.. - Rohanimanesh, Mahadevan - 2001
- 5 Continuous-time hierarchical reinforcement learning - Ghavamzadeh, Mahadevan - 2001
- 4 Integrating experimentation and guidance in relational reinf.. - Driessens, Dzeroski - 2002
- 3 Lyapunov design for safe reinforcement learning - Perkins, Barto
- 3 Lyapunov-constrained action sets for reinforcement learning - Perkins, Barto - 2001
- 2 Statecharts: A visual formalism for complex systems (context) - Harel - 1987
- 2 Learning hierarchical control structure from multiple tasks .. (context) - Digney - 1998
- 2 ce-navigation robot (context) - Nourbakhsh, Powers et al. - 1995
- 2 Singular Perturbation Methodology in Control Systems (context) - Naidu - 1988
- 2 Structured approximation of stochastic temporally extended a.. (context) - Rohanimanesh, Mahadevan
- 2 Automated state abstraction for options using the U-tree alg.. - Jonsson, Barto - 2000



The graph only includes citing articles where the year of publication is known.

Documents on the same site (<http://www.cs.umass.edu/~mahadeva/organized-pubs.html>): [More](#)
 A Multiagent Reinforcement Learning Algorithm by.. - Ghavamzadeh, Mahadevan (2002) ([Correct](#))
 Approximate Planning with Hierarchical Partially.. - Theodorou, Mahadevan (2002) ([Correct](#))
 Extending Hierarchical Reinforcement Learning to.. - Ghavamzadeh, Mahadevan ([Correct](#))

[Online articles have much greater impact](#) [More about CiteSeer.IST](#) [Add search form to your site](#) [Submit documents](#) [Feedback](#)

CiteSeer.IST - Copyright [Penn State](#) and [NEC](#)